

Date of Deposit 06/21/00

I hereby certify that this paper or fee is being deposited with the United States Postal Service "Express Mail Post Office to Addressee" service under 37 CFR 1.10 on the date indicated above and is addressed to the Commissioner of Patents and Trademarks, Washington, D.C., 20231.

Mailers: LORNA L. SCHOLZ

Malle (Signature)

TRW Docket No. 22-0124

- 1 -

Beam Hopping Self Addressed Packet Switched

Communication System with Power Gating

5

CROSS-REFERENCE TO RELATED APPLICATIONS

This application is related to TRW Docket No. 22-0132, titled "Beam Hopping Self Addressed Packet Switched Communication System with Multi-port Memory, filed ___, as serial No. ____; TRW Docket No. 22-0131, titled "Beam Hopping Self Addressed Packet Switched Communication System with Locally Intelligent Scheduling", filed ___, as serial No. ____; TRW Docket No. 22-0133, titled "Beam Hopping Self Addressed Packet Switched Communication System with Multiple Beam Array Antenna", filed ___, as serial No. ____; TRW Docket No. 22-0006, titled "Gated Power Time Division Downlink for a Processing Satellite", filed March 16, 1999 as Serial No. 09/270,361; and TRW Docket No. 22-0127, title "Beam Hopped Gated Power

DISCOUNT

- 2 -

Time Division Downlink", filed ____ as Serial No. ____.

5

BACKGROUND OF THE INVENTION

The present invention relates to satellite communication systems. In particular, the present invention relates to a beam hopping satellite communication system including a self addressed packet switch that power gates a downlink beam.

Satellites have long been used to provide communications capabilities on a global scale. Since the inception of the modern communications satellite, however, one factor has remained constant: the limited availability of power on board the satellite. The limited availability of power persists today even in the face of tremendous advances in satellite technology.

- 3 -

The communications reception equipment used to receive the uplink and the transmission equipment used to generate the downlink are significant drains on satellite power. The transmission equipment in particular often requires 50% or more of the total power generated by a satellite. Furthermore, downlink power amplifiers are far from 100% efficient and therefore waste power whenever they are active.

Any undue drain on satellite power is undesirable. Thus, for instance, limitations on satellite power may prevent a satellite from encoding and decoding heavier and more robust coding techniques. As another example, limited satellite power may reduce the number and type of observational or sensing functions which a satellite may perform, or the number of downlinks and uplinks a satellite may support.

In addition, as satellite technology has progressed, it has become more common for satellites to process their uplinks. In other words, the satellite may decode, process, route, queue, and

- 4 -

otherwise manipulate data before recoding and packaging the data into downlink frames. Thus, the utilization of the downlink depends on the amount of data ready and waiting to be transmitted.

5 Transmitting partially empty frames can be a waste of power, and can have detrimental impacts on the satellite performance through unnecessary servicing of queues, for example.

10 A need has long existed in the industry for a communication system that addresses the problems noted above and others previously experienced.

BRIEF SUMMARY OF THE INVENTION

15 A preferred embodiment of the present invention provides a downlink beam frame signal processing system for a communication satellite. The processing system includes a packet switch that routes self addressed uplink data to a memory. The memory stores queues for at least a first and a second downlink beam
20 hop locations. The processing system also includes a

- 5 -

power amplifier that amplifies a waveform formed using the uplink data. A power gating circuit coupled to the power amplifier includes a power gate input responsive to a power gating signal to remove power from at least a portion of the waveform before transmission.

The power gating signal may be asserted in response to unavailability of downlink data in the memory. As examples, the power gating signal may be asserted when there is not enough data in the memory to fill one or more payload fields in the waveform, in order to meet satellite power requirements, in response to statistical multiplexing estimates, or to meet desired average queue depths.

Another preferred embodiment of the present invention provides a method for processing a downlink beam frame signal. The method includes the step of switching self addressed uplink data into at least one of a first and second downlink hop location storage area in a memory, amplifying a waveform based in part on the uplink data for transmission, and prior to

- 7 -

Figure 5 depicts the odd hop downlink beams from the beam laydown of Figure 3.

Figure 6 shows an implementation of a router.

Figure 7 shows an implementation of an inbound
5 module.

Figure 8 illustrates an implementation of an outbound module.

Figure 9 shows a cell discard algorithm for fixed partition buffers.

10 Figure 10 shows a cell discard algorithm for dynamically buffered queues.

Figure 11 illustrates a method for routing data through a satellite to a selected downlink hop location.

15 Figure 12 shows a modulator implementation that supports power gating.

- 8 -

Figure 13 shows a multiple payload frame signal with exemplary power gating control signals.

Figure 14 illustrates operational steps that occur before and after power gating a beam hopping multiple payload frame signal.

Figure 15 shows a preferred embodiment of a downlink frame suitable for power gating.

DETAILED DESCRIPTION OF THE INVENTION

Turning now to Figure 1, that figure shows a block diagram of a downlink beam processing system or bandwidth switch 100. The bandwidth switch 100 includes a controller 102 and a waveform processing chain that operates on data provided by the data source 104. In particular, the waveform processing chain includes a waveform generator 106, an amplifier 108, and a feed switch 110. The waveform processing chain further includes a first feed path 112 and a second feed path 114 that may be characterized by a

- 9 -

polarization effect on the waveform that propagates along the feed paths 112-114. The polarization effect may induce, for example, clockwise (right) or counter clockwise (left) polarization in the waveform.

5 The first feed path 112 terminates in a first radiating element 116 (e.g., a feed horn). Similarly, the second feed path terminates in a second radiating element 118 (e.g., another feed horn). The first and second feed horns 116, 118 illuminate the subreflector
10 120. The subreflector 120, in turn, illuminates the main reflector 122 that projects downlink beams onto terrestrial cells. Thus, the first and second feed horns 116, 118, the subreflector 120, and the main reflector 122 form a Multiple Beam Array Antenna (MBA)
15 to direct spot beam coverage to distinct terrestrial cells. Additional feed horns may be used with the MBA to generate additional spot beams, and multiple independent MBAs may be provided.

 The waveform generator 106 accepts baseband data
20 from the data source 104 and creates a waveform to be

- 10 -

transmitted (after amplification by the amplifier 108). The switch 110 selects the particular feed path 112-114 along which the waveform propagates (and thus, in certain embodiments, the polarization and/or hop location associated with the waveform).

The controller 102 exercises color control over the waveform to be transmitted. Thus, the controller 102 may output one or more control signals (collectively referred to as a color selection signal) that determine, for example, the frequency, polarization, or hop location of the waveform to be transmitted. In the preferred embodiment, the beam color components include Even and Odd hop locations, Left and Right polarization, and first and second frequencies. Eight different colors are therefore available: 1EL, 1ER, 1OL, 1OR, 2EL, 2ER, 2OL, 2OR.

With regard to Figure 2, a more specific implementation of a downlink beam processing system and bandwidth switch 200 is shown. The bandwidth switch 200 includes a data scheduler 202, a data

- 11 -

router 204, and a waveform processing chain including a QPSK modulator 206, an upconverter 208, and a traveling wave tube amplifier (TWTA) 210. The switch 110 is illustrated in Figure 2 as a ferrite switch 110 that directs the waveform to be transmitted through either the first feed path 112 or the second feed path 114. Preferably, additional ferrite switches 212 and 214 in the feed paths 112-114 provide additional signal isolation (e.g., approximately 20db between input and output when the ferrite switch is off). The additional ferrite switches 212-214 operate under control of the color selection output to pass or block a waveform to be transmitted through the feed paths 112-114. In other words, when the waveform to be transmitted is destined for the feed 112, then the ferrite switch 214 is coupled through the load 228 to ground. Similarly, when the waveform to be transmitted is destined for the feed 114, then the ferrite switch 212 is coupled through the load 226 to ground.

- 12 -

In addition, Figure 2 shows a color selection output 216, two frequency selection inputs 218 and 220, a feed path selection input 222, and an intermediate waveform output 224.

5 During operation, the bandwidth switch 200 accepts baseband data from the router 204 (e.g., an ATM cell router), and creates a waveform to be transmitted using the waveform processing chain. The waveform processing starts by directly converting
10 baseband I and Q data to an intermediate frequency of, for example, 750 MHz. The waveform processing then selects one of F1 (e.g., 3.175 MHz) and F2 (e.g., 3.425) and one of F3 (e.g., 16 GHz) and F4 (e.g., 17.4 GHz) to produce a waveform to be transmitted with a
15 final center frequency at one of 18.425 GHz, 18.675 GHz, 19.825 GHz, and 20.075 GHz. The scheduler 204 monitors the propagation of data through the waveform processing chain and determines the color of the waveform to be transmitted. To that end, the
20 scheduler 204 provides the color selection output 216

Case	Age	Sex	Occupation	Duration of illness	Site of lesion	Pathological changes	Microscopic findings	Diagnosis
1	25	M	Farmer	10 years	Brain	Chronic	Microscopic	Chronic
2	35	F	Housewife	5 years	Brain	Chronic	Microscopic	Chronic
3	45	M	Teacher	15 years	Brain	Chronic	Microscopic	Chronic
4	55	F	Retired	20 years	Brain	Chronic	Microscopic	Chronic
5	65	M	Engineer	25 years	Brain	Chronic	Microscopic	Chronic
6	75	F	Homemaker	30 years	Brain	Chronic	Microscopic	Chronic
7	85	M	Retired	35 years	Brain	Chronic	Microscopic	Chronic
8	95	F	Retired	40 years	Brain	Chronic	Microscopic	Chronic

Case	Age	Sex	Occupation	Duration of illness	Site of lesion	Pathological changes	Microscopic findings	Diagnosis
1	25	M	Farmer	10 years	Brain	Chronic	Microscopic	Chronic
2	35	F	Housewife	5 years	Brain	Chronic	Microscopic	Chronic
3	45	M	Teacher	15 years	Brain	Chronic	Microscopic	Chronic
4	55	F	Retired	20 years	Brain	Chronic	Microscopic	Chronic
5	65	M	Engineer	25 years	Brain	Chronic	Microscopic	Chronic
6	75	F	Homemaker	30 years	Brain	Chronic	Microscopic	Chronic
7	85	M	Retired	35 years	Brain	Chronic	Microscopic	Chronic
8	95	F	Retired	40 years	Brain	Chronic	Microscopic	Chronic

Case	Age	Sex	Occupation	Duration of illness	Site of lesion	Pathological changes	Microscopic findings	Diagnosis
1	25	M	Farmer	10 years	Brain	Chronic	Microscopic	Chronic
2	35	F	Housewife	5 years	Brain	Chronic	Microscopic	Chronic
3	45	M	Teacher	15 years	Brain	Chronic	Microscopic	Chronic
4	55	F	Retired	20 years	Brain	Chronic	Microscopic	Chronic
5	65	M	Engineer	25 years	Brain	Chronic	Microscopic	Chronic
6	75	F	Homemaker	30 years	Brain	Chronic	Microscopic	Chronic
7	85	M	Retired	35 years	Brain	Chronic	Microscopic	Chronic
8	95	F	Retired	40 years	Brain	Chronic	Microscopic	Chronic

- 14 -

In addition, either of the feed paths 112-114 may be characterized by a polarization effect on the waveform that propagates along the feed path. Thus, the color selection output 216 may also determine the polarization color component of the waveform to be transmitted. Optionally, however, separate feed paths may be provided for any number of desired combinations of polarization and hop location. The transmitted waveform manifests itself as a beam spot that, typically, provides downlink bandwidth for a terrestrial cell.

The bandwidth switch 200 may operate onboard a first satellite that supports a cellular coverage area using a set of spot beams. The scheduler 202 ensures that the waveforms to be transmitted have the appropriate beam colors to minimize co-channel, adjacent channel, and cross polarization for the cellular coverage area and the eight possible beam colors. However, when, for example, a second subsequently launched satellite begins to provide

- 15 -

bandwidth support for the same cellular coverage area, the bandwidth switch 200 allows the first satellite to modify its beam colors to accommodate the second satellite. In other words, the bandwidth switch 200
5 allows the first and second assignment of spot beams to the coverage area to coexist in a minimally interfering manner. The resultant beam laydown may then be minimally interfering initially for a single satellite, and later reconfigured to be minimally
10 interfering with regard to a particular type of interference or interferences for additional satellites providing bandwidth for the same coverage area.

Turning next to Figure 3, that figure illustrates
15 a beam laydown 300 that uses hopping beams. The coverage area is generally divided into cells as shown in idealized form, for example, by the hexagonal cells 302 and 304. Each of the cells is also labeled with a beam color. For example, a beam of color 10L provides

- 16 -

bandwidth for the cell 302, while a beam of color 2EL provides bandwidth for the cell 304.

The laydown 300 is characterized in that, for mutually exclusive hop locations, a only six co-
5 channel interferers (CCI) (caused by a beam of the same color), zero adjacent channel interferers (ACI) (caused by a beam differing in only one color component), and zero cross polarization interferers (XPI) (caused by a beam differing only in
10 polarization) exist for any given cell. In other words, taking cell 306 (color 1ER) as an example, the CCIs are cells 308, 310, 312, 314, 316, and 318.

Note that cell 320 does not provide CCI because it has an odd color component and is not provided with
15 spot beam energy at the same time as the cell 306 (color 1ER) (i.e., the hop locations are mutually exclusive). The laydown 300 also provides minimal interference when hop locations are non-mutually exclusive. In the non-mutually exclusive case, there
20 exist only 6 CCIs, 2 ACIs, and 2 XPIs. The ACIs are

- 17 -

cells 322 and 324, while the XPIs are cells 320 and 326. Note that not all colors (e.g., 20L) need be used in a beam hopping beam laydown.

Figure 4 shows the laydown 300 as well. In Figure 4, however, only the even hop locations are marked. Similarly, Figure 5 shows the beam laydown 300 with only the odd hop locations marked.

Turning next to Figure 6, a preferred implementation of a router 600 is illustrated. The router 600 includes thirty-five inbound modules (IBMs), three of which are designated 602, 604, 606. The IBMs 602-606 are coupled to input ports of an ATM cell switch 608. The ATM cell switch 608 has thirty-three outputs coupled to individual outbound modules (OBMs), three of which are designated 610, 612, 614. pairs of uplink demodulators feed each IBM 602-606, while the OBMs 610-614 feed downlink modulators.

The router 600 provides a self addressed packet switching function. In other words, the router 600 uses addressing or destination information present in

- 18 -

uplink data (e.g., ATM cells) to deliver the cells to a specific data queue that feeds a downlink beam appropriate for the destination or next hop of the cell. Thus, for example, the VPI / VCI fields in an ATM cell may be used to guide the cell into an appropriate downlink queue. Cells may first be discarded however, if they fail their header error check.

The output of the IBMs 602-606 includes a routing tag, a queue tag, and the (possibly modified) cell itself. The role of the IBMs 602-606, the routing tag and, the queue tag will be described in more detail below with respect to Figure 7. In general, the ATM cell switch 608 uses the bits in the routing tag to connect a cell switch input port to a cell switch output port. The queue tag, a portion of the routing tag, and the cell itself then flow through the switch to the OBM connected to the selected output port. As will be described in more detail below, each OBM 610-614 includes a set of downlink queues that feed

- 19 -

downlink beams directed to predetermined terrestrial cells. The queue tag determines in which downlink queue the cell will be inserted in the OBM (and may be indicative of cell priority and downlink coding rate).

5 Thus, the IBMs 602-606, the ATM cell switch 608, and the OBMs 610-614 operate in concert to deliver cells to an appropriate downlink queue in a self addressed manner.

Figure 7 illustrates an implementation 700 of the
10 IBMs 602-606. In particular, the implementation 700 includes a lookup or routing table 702 and an output buffer 704. An incoming ATM cell generally indicated at 706 is shown to include (among other fields) a payload 708 and a VPI/VCI address 710. Figure 7 also
15 illustrates a particular routing tag 712, queue tag 714, and optional replacement VPI/VCI 716 field for the cell from among those stored in the routing table 702. If the cell is modified (e.g., by changing its VPI/VCI), the IBM will also recompute the cell header
20 error check. A ground based Network Control Center

- 20 -

(NCC) may dynamically update the routing table 702 to ensure proper routing of cells from the current network node (e.g., the satellite) to the next network node (e.g., a ground terminal).

5 The VPI/VCI 710 of the cell 706 addresses the routing table 702. In response, the routing table 702 provides the routing tag 712, queue tag 714, and new VPI/VCI addresses 716 (e.g., for the next hop that the cell will make). A NULL entry in the routing table
10 702 may indicate that the cell is to be discarded. Any modifications to the uplink cell result in the IBM recomputing an error check for the uplink cell as well. This information enters the output buffer 704 (which may be, for example, 8191 cells in length).
15 Once in the output buffer 704, the information awaits selection by an arbitration algorithm before it enters the cell switch 608. As examples, the arbitration algorithm may give preference to the oldest cells (e.g., using a two bit quantization of clock cycle
20 cell age), the remaining capacity of the output buffer

- 21 -

704 (e.g., using a three bit quantization of total input queue size), and the like.

Once the cell is selected, its routing tag is used to send the cell to an associated output port of the cell switch 608. In particular, the routing tag 712 is preferably seven bits in length. The ATM cell switch 608 uses six of the seven bits internally to connect an input port to an output port determined by the six bits. For future expandability, the seventh bit may be used, for example, to support larger switches with additional output ports.

Turning now to Figure 8, that figure illustrates an implementation of an OBM 800. The OBM 800 includes an OBM controller 802 coupled to an external cell memory 804. The OBM controller 802 integrates, preferably, into a single ASIC logic that implements a switch interface controller (SIC) 806 and a switch interface data handler (SID) 808. The SIC 806 couples to a downlink schedule table 810, a queue statistics memory 812, a linked list memory 814, and a pointer

- 22 -

memory 816. In addition, the OBM 800 includes a first Reed-Solomon encoder (RSE) 818, a second RSE 820, interface electronics (IEA) 822 coupled to interleaving memory 824, and a downlink frame
5 formatter (DLF) 826.

The external cell memory 804 is preferably organized into numerous queues. The queues may be distinguished by characteristics such as hop location, priority, and code rate, or other criteria. In
10 general, for each hop location, there may be one or more code rates, each having one or more priority queues. In one embodiment, there are 16 downlink hop locations (referred to as a subclass) the external cell memory 804 includes 16 light coding queues and 16
15 heavy coding queues (i.e., 512 total queues). Each of the 16 light and 16 heavy coding queues represents a predetermined priority. One queue (e.g., priority 15, subclass 15, light coding) may be reserved for system controller traffic. The queue tag determines the
20 subclass and the queue for which a cell is destined.

- 23 -

The external cell memory 804 is preferably a multiport memory shared between output ports of the cell switch 608. The multiport nature of the external cell memory 804 resides in its role as shared storage for multiple
5 hop locations (i.e., Beam A and Beam B that share a single physical cell switch 608 output port) served by the single OBM controller 802.

The memory provided by the external cell memory 804 may be allocated in a fixed or dynamic manner in
10 several different ways. As one example, one or more queues may be allocated a fixed amount of memory to meet the expected long term needs of the subclass and priority associated with the queue. The remaining memory may then be shared by the remaining queues. To
15 guarantee a minimum bandwidth for each queue, a minimum threshold amount of memory may be reserved for each queue. Thus, the external cell memory 804 permits pairing destination bandwidth needs for a particular destination at a particular time with
20 allocations of queue memory. To that end, the NCC may

- 24 -

dynamically uplink to the satellite changes to the manner in which memory is allocated. The thresholds, maximum queue size, minimum queue size, and the like are stored in the pointer memory 816.

5 The SIC 806 comprises logic that directs the activities of the OBM controller 802, including obtaining cells from the cell switch 608 through the SID 808. As will be described in more detail below, the SIC 806 makes a determination regarding whether
10 the cell should be accepted or rejected using parameters for each queue stored in the pointer memory 816. If a cell is accepted, the SID 808 stores the cell in a queue in the external cell memory 804. The SIC 806 then updates the linked list memory 814 to
15 record where the cell was stored in the external cell memory 804. The SIC 806 also updates the queue statistics memory 812 to reflect the number of cells in each queue in the external memory 804, the number of cells accepted or rejected for each queue, peak
20 queue occupancies, the number of cells pulled from the



5

10

15

20

- 26 -

convolutional code for heavy coded cells. The DLF 826 then forms, preferably, a two payload downlink frame for the downlink, including overhead information (e.g., synchronization codes, code identifiers, guard
5 time, and the like). Each payload may independently carry 12 heavy coded cells or 24 light coded cells.

Additional details of the frame format, coding, interleaving, and scrambling may be found in TRW Docket No. 22-0125, Serial No. ____, incorporated
10 herein by reference in its entirety.

As noted above, the SIC 806 makes a determination regarding whether a cell retrieved from the cell switch 608 should be accepted or rejected using parameters for each queue stored in the pointer memory
15 816. Turning now to Figure 9, a flow diagram 900 presents a series of determinations made by the SIC 806 when queue are fixed in size. At step 902, the SIC 806 determines whether there is any free memory in which to store the cell. If the free cell counter
20 (FCC, i.e., the total buffer size minus the number of

- 27 -

queued cells) is zero, the cell is discarded (step 904). Otherwise the SIC 806 determines if the cell is a controller cell (step 906). A controller cell, for example, may be a cell that carries command, configuration, or status information from the NCC to the satellite (e.g., to update the routing table 702 or downlink scheduling table). If the cell is a controller cell, and if the associated queue depth (QD) is less than its maximum size (i.e., the All_Thr), then the cell is accepted, otherwise it is discarded (step 908).

Continuing at step 910, if the queue depth is less than or equal to the minimum threshold queue size (Min_Thr) then the cell is accepted (step 912). Step 914 checks to see if the queue depth is greater than the maximum allowed queue size (Max_Thr), and if so the cell is discarded (step 916). Beginning at step 918 the SIC 806 may accept or discard the cell based on the Cell Loss Priority (CLP) field found, for example, in an ATM Cell.

- 28 -

A CLP of zero indicates that the cell is of high priority and should not be dropped during periods of high congestion. A CLP of one indicates that the cell is of low priority and should be dropped, if needed, during periods of high congestion. At step 918, if the cell is low priority, and the queue depth is greater than the cell loss priority threshold (CLP_Thr), then the cell is discarded (step 920). If (step 922) the queue depth is greater than All_Thr, then the cell is discarded (step 924). Otherwise, the cell is accepted (step 926).

When queues are allocated memory in a dynamic fashion, the SIC 806 follows the steps indicated in Figure 10 to determine whether a cell is to be accepted. In particular, at step 1002, the SIC 806 determines whether there is any free memory available to store the cell. If the free cell counter (FCC) is zero, the cell is discarded (step 1004). Otherwise the SIC 806 determines if the cell is a controller cell (step 1006). If so, if the queue depth (QD) is

- 29 -

less than All_Thr, then the cell is accepted,
otherwise discarded (step 1008).

Continuing at step 1010, if the queue depth is
less than or equal to the minimum threshold queue size
5 (Min_Thr) then the cell is accepted (step 1012). Step
1014 checks to see if the queue depth is greater than
the maximum allowed queue size (Max_Thr), and if so
the cell is discarded (step 1016). At step 1018 the
SIC 806 may accept or discard the cell based on the
10 CLP. If the cell is low priority, and the amount of
free memory is less than or equal to the cell loss
priority threshold (CLP_Thr), then the cell is
discarded (step 1020). If (step 1022) the amount of
free memory is less than All_Thr, then the cell is
15 discarded (step 1024). Otherwise, the cell is
accepted (step 1026).

The pointer memory 816 stores the thresholds
referred to above, including the All_Thr, Min_Thr,
Max_Thr, CLP_Thr, and FCC for each queue.

- 30 -

Returning again to the external cell memory 804, its shared nature allows a queue to grow in size to handle bursts of traffic. In addition, dynamic buffers allow the NCC to accept connections whenever shared memory is available to allow a queue to grow. On the other hand, a fixed partition buffer must be checked to determine if it has any room left for cells generated by the new connection. Bandwidth available only at a different priority level nevertheless denies the connection. The shared memory approach, however, allows the appropriate priority level queue to grow in size to handle the connection. Network and queue management functions therefore tend to be less complex and more efficient.

It is also noted that the external cell memory 804 may also include a queue dedicated to controller cells. The downlink scheduling table services the controller cells by placing them in a controller buffer (as opposed to preparing them for transmission in a downlink frame), for example, 32 or 64 cells in

- 31 -

size. One or more control elements may then access the controller buffer, decode commands in the cell, and perform those commands. As an example, a controller cell may direct an inbound module to
5 replace entries in the routing table 702.

Figure 11 summarizes a method 1100 for routing data through a satellite to a selected downlink hop location. At step 1102, the satellite looks up a hop location destination queue using an address carried in
10 the uplink data. Next, at step 1104, the satellite switches the uplink data through a switch, and stores (step 1106) the data in the appropriate queue.

In building frames for transmission, the satellite, at step 1108, first retrieves data from the
15 queue in order to build the downlink waveform. The satellite then selects a feed path for the waveform according to its destination hop location (step 1110). The waveform is transmitted (step 1112), preferably using a multiple beam array antenna with feed elements
20 assigned to the hop locations.

- 32 -

As noted above, the downlink beam may be power gated in a number of ways to save power, or to achieve other goals. Turning next to Figure 12, that figure shows an implementation 1200 of the modulator 206 that supports power gating. Inphase data is supplied to the Inphase gate 1202 while Quadrature data is supplied to the Quadrature gate 1204. As illustrated, the Inphase and Quadrature gates 1202, 1204 are D flip flops with reset inputs. The Inphase and Quadrature gates 1202, 1204 feed a digital modulator core 1206 that produces a modulated waveform on a modulator output 1208. A local oscillator (LO) signal (preferably 750 MHz) provides an intermediate frequency carrier signal. The amplifier 1210 boosts the modulated waveform, after which it is filtered by the bandpass filter 1212. The bandpass filter 1212 preferably has a passband centered at 750 MHz, for example, from 625 to 875 MHz.

A data clock 1214 that preferably runs at 196.7 MHz drives the Inphase and Quadrature gates 1204,

- 33 -

1204. Note that a power gate input 316 connects to the Inphase and Quadrature gates 1202, 1204, as well as to the gating control input 1218 of the digital modulator core 1206. When an active power gating
5 signal is present on the power gate input 1216, the Inphase and Quadrature gates 1202, 1204 have their outputs held in a known state (e.g., both 0). Furthermore, the digital modulator core 1206 outputs a signal with frequency content outside of the passband
10 of the bandpass filter 1212.

For example, the digital modulator core 1206 may output a DC signal in response to the active power gating signal. As a result, the bandpass filter eliminates the DC signal. A power gated signal
15 results.

Returning to Figure 2, the upconverter 208 (e.g, a 20 GHz mixer) ordinarily outputs a fully upconverted signal for amplification and transmission. However, the absence of energy in the power gated signal causes
20 the upconverter to produce substantially no signal at

- 34 -

its output during power gating. As a result, the TWTA 210 does not expend amplification energy, and substantially no downlink energy is present in the downlink beam while the power gating signal is active.

5 In other words, the DC power consumption of the TWTA 210 is reduced by substantially eliminating radiated power.

Turning next to Figure 13, that figure present a timing diagram 1300 that illustrates a multiple
10 payload frame signal 1302 and power gating signals 1304, 1306, 1308, 1310, 1312, 1314, 1316 (assumed active when high). As an example, the frame signal 1302 may include a 368 symbol first header signal 1318, a 7552 symbol first payload signal 1320, a first
15 16 symbol flush signal 1322, a 96 symbol second header signal 1324, a 7552 symbol second payload signal 1326, and a second 16 symbol flush signal 1328. In general, however, the frame signal 1302 may include N headers and N payloads independently subject to power gating.

- 35 -

The power gating signal 1304 never goes active during the frame signal 1302. Thus, none of the frame signals 1318-1328 are power gated. As a result, both first and second header signals 1318, 1324 and both
5 first and second payload signals 1320, 1326 are delivered to the ground. In contrast note that the power gating signal 1316 is active across the entire frame signal 1302. Thus, substantially no energy is provided in the downlink beam over the time during
10 which the frame signal 1302 would be transmitted.

On the other hand, the power gating signal 1306 goes active during the second payload signal 1326 and the second flush signal 1328. Thus, the frame signal 1302 continues to bear important overhead information
15 in the first and second header signals 1318, 1324. The overhead information may include, for example, synchronization bits, beam hopping location identifiers, frame coding identifiers, frame counts, and the like.

- 36 -

The overhead information may further include power gating bit patterns that indicate to a ground receiver which frame signals are power gated. As an example, the first header signal 1318 or second header
5 signal 1324 may include a frame type field that carries repetitions of the bit pattern 10100101 to indicate power gating of the first payload signal 1320 or second payload signal 1326, or repetitions of the bit pattern 11110000 to indicate power gating of the
10 entire frame signal 1302. In particular, bit patterns may be assigned to identify any combination of header, payload, and flush signal power gating. Note also that a ground receiver may deactivate its own receivers in response to the bit patterns to save
15 power during power gated sections of the frame signal.

Still with reference to Figure 13, the power gating signal 1310 results in power gating of the first and second payload and flush signals 1320, 1322, 1326, 1328. Similarly, the power gating signal 1312

- 37 -

results in power gating of the first payload signal 420 and the first flush signal 1322.

Because the multiple payload frame signal 1302 includes multiple headers, each preferably bearing
5 synchronization information, additional power gating options are available. Thus, for example, the power gating signal 1308 power gates the second header signal 1324, second payload signal 1326, and the second flush symbols 1328. Synchronization is
10 nevertheless provided by the first header signal 1318. Similarly, the power gating signal 1314 power gates all the frame signals except for the first header signal 1318.

The scheduler 202 may include logic to assert the
15 power gating signal under many scenarios. For example, when the satellite moves into eclipse and less power is available, the scheduler 202 may power gate every Nth complete frame, every other payload, or any combination of frame signals to achieve a desired
20 power reduction. As another example, the scheduler

- 38 -

202 may activate the power gating signal in response to a statistical multiplexing estimate of downlink beam utilization. As an example, if the downlink beam is estimated to be 90% utilized during a certain time period, then the scheduler 202 may power gate 1-10% of the frames or payloads. Such estimates may be uplinked to the satellite or generated onboard.

As another example, the scheduler 202 may determine when to activate power gating based on the current terrestrial cell hop location of the downlink beam. Thus, scheduler 202 may power gate the second payload signal 1326 if the bandwidth requirements of the current terrestrial cell are met by the first payload signal 1320 alone. As yet another example, the scheduler 202 may power gate based on data queues present in the router 204. For example, a data queue from which ATM data cells are extracted to fill the second payload signal 426 may consistently have too few cells to completely fill the second payload signal 1326. In response, the scheduler 202 may power gate

- 39 -

the second payload signal 1326 periodically in order to maintain the data queue approximately at preselected occupancy level, on average.

Yet another example of conditions that give rise to power gating include lack of cells in the data queues. For example, if there are no cells in the data queues that feed a downlink beam, then the downlink frame may be power gated until cells are available. In addition, if too few cells are available to fill the first payload, then the first payload may be power gated. Similarly, if too few cells are available to fill the second payload, then the second payload may be power gated. Alternatively, the entire frame may be power gated until enough cells are available to fill both payloads of a frame. To this end, power gating thresholds set by the NCC may be uplinked to the satellite. The power gating thresholds indicate how many cells need be available for each payload, or the frame in general, before the payload is transmitted without power gating. The

- 40 -

remaining space in the payload may then be filled, for example, with NULL cells or other placeholder data.

Turning next to Figure 14, that figure shows a flow diagram 1400 of the operational steps that occur before and after power gating. The operational steps include hopping a downlink beam between at least two terrestrial cells (step 1402). At step 1404, queue depths are monitored and statistical multiplexing estimates of downlink utilization are obtained. At step 1406, power gating is activated based on, as examples, beam hop locations, power saving goals, queue depths, the number of cells available in the queues, and the like, as discussed in detail above.

Continuing at step 1408, one or more header signals, payload signals, and flush signals are power gated. Thus, at step 1410, a frame signal is transmitted in which at least one header signal, payload signal, or flush signal has substantially no energy in the downlink beam.

- 41 -

Figure 15 shows a preferred embodiment of a downlink frame signal 1500 suitable for power gating as discussed above. The frame signal 1500 includes a first header field 1502 followed by a first payload field 1504 and a first flush field 1506. In addition, the frame signal 1500 includes a second header field 1508 followed by a second payload field 1510 and another flush field 1512. The first header field 1502, first payload field 1504, first flush field 1506, second header field 1508, second payload field 1510, and second flush field 1512 are all encapsulated into the single frame 1500.

Continuing with reference to Figure 15, the first header field 1502 is composed of several subfields. In particular, the first header field 1502 includes a hopping beam guard band 1514, a first payload pseudorandom noise (PN) synchronization field 1516, and a spare field 1518. The first header field 1502 also includes a first frame type field 1520, a

- 42 -

masterframe count field 1522, and a subframe count field 1524.

The second header section includes a smaller set of subfields, namely, the second PN synchronization field 1526 and the second frame type field 1528.

Table 1, below, shows the preferred length and modulation of each field. Symbols are preferably transmitted at 196.7 megasymbols per second.

Table 1		
Field	Symbols	Modulation
first header 1502	368	
hopping beam guard band 1514	114	BPSK
first payload PN synch 1516	64	BPSK
spare 1518	62	BPSK
first frame type 1520	32	BPSK
masterframe count 1522	32	BPSK
subframe count 1524	64	BPSK
first payload 1504	7552	QPSK
first flush 1506	16	QPSK
second header 1508	96	
second payload PN synch 1526	64	BPSK
second frame type 1528	32	BPSK
second payload 1510	7552	QPSK

5
10

15

The first PN synchronization field 1516 and the second PN synchronization field provide synchronization bits for earth terminals. As will be explained in more detail below, a single PN synchronization sequence generator is used to provide an identical PN sequence for both PN synchronization fields 1516, 1528. The subframe count field 1524 counts individual frames as they are transmitted. Preferably, the subframe count field 1524 includes a

- 44 -

16 bit downlink frame count appended with 8 zeros and convolutionally encoded with a relatively heavy (e.g., 3/8 rate) code. The masterframe count field 1522 increments at the start of every masterframe (e.g., every 9328 frames). The masterframe count rolls over after reaching its maximum value (0xFFFFFFFF), although it may be reset or preprogrammed at any time.

The spare field 1518 may be drawn from to provide subsequent enhancements to the frame 1500 (e.g., additional synchronization bits). Preferably, the spare field 1518, the hopping beam guard band 1514, and first PN synchronization field 1516 are filled with PN bits that are generated by a PN synchronization sequence generator discussed below.

The first frame type field 1520 generally indicates characteristics of the first payload field 1504, while the second frame type field 1528 generally indicates characteristics of the second payload field 1510. Several examples of codes for the first and

- 45 -

second frame type fields 1520, 1528 are illustrated below in Table 2.

Table 2		
Frame Type	Uncoded Value	Coded Value
Light Coding	110	00111100
Heavy Coding	011	10010110
Frame Gate	001	10100101
Power Gate	000	11110000

where the coded values may result from the application of an (8, 4) Reed-Muller block code.

Although the light coding, heavy coding, and power gating options are with reference to a payload itself, the frame gate option indicates power gating of an entire frame (i.e., all 15600 symbols). Each coded value is preferably repeated four times in the frame type field. For example, a frame type of 00111100 00111100 00111100 00111100 in the first frame type field 1520 indicates that the first payload field

- 46 -

104 is lightly coded. As another example, a frame type of 11110000 11110000 11110000 11110000 in the second frame type field 1528 indicates that the second payload field 1510 will be power gated. When a frame or payload field is power gated, only a small fraction of the ordinary output power will be generated in the downlink beam during for the entire frame, or during the identified payload(s).

With regard to the heavy coding and light coding, as examples, a lightly coded payload may indicate 3/4 rate, constraint length 7, punctured convolutional coding of 1416 Reed-Solomon block coded bytes. A heavily coded payload may indicate a 3/8 rate, constraint length 7, punctured convolutional coding of 708 Reed-Solomon block coded bytes. Thus, the first and second payload fields remain the same size (7552 symbols) under both coding rates.

The first and second payload fields 1504, 1510 carry the billable data to the earth terminals. The first and second payload fields 1504, 1510 are

- 47 -

typically concatenated coded using an inner convolutional code. The, the first and second flush fields 1506, 1512 are provided as a convenient way for the earth terminal convolutional decoders to reset
5 their state in preparation for the next payload.

The frame signal 1500 delivers multiple payloads (in the preferred embodiment, two payloads) in a single frame. Although a first header field 1502 is provided as well as a second header field 1508, the
10 second header field 1508 is smaller than the first header field 1502. In particular, the second header field does not repeat the hopping beam guard band 1514 (since the receiver(s) for the first and second payload fields 1504, 1510 are in the same beam spot
15 for the current hop location), spare field 1518, masterframe count 1522 and subframe count 1524 (since only one count is needed for the single multiple payload frame).

As a result, the frame 1500 delivers two payloads
20 in a single frame with less overhead than would be

- 48 -

incurred by transmitting two single payload frames. Throughput is therefore higher. The specific frame 1500 shown in Figure 1 may be generalized to a single N payload N header frame, under the general condition
5 that the sum of the overhead arising from the N headers is less than the sum of the overhead arising from N individual single payload frames.

Thus, the present invention provides a beam hopping self addressed packet switched communication
10 system with power gating. In particular, the communication system reduces power consumption through power gating downlink beam frames in several situations. As examples, the communication system may power gate when there is no data to send in the
15 downlink beam, where there is not enough data to build a complete payload, or when there is not enough data to build a complete frame. As additional examples, the communication system may power gate to achieve a desired average queue depth, or according to a

- 49 -

statistical multiplexing estimate of downlink utilization.

While the invention has been described with reference to a preferred embodiment, those skilled in the art will understand that various changes may be made and equivalents may be substituted without departing from the scope of the invention. In addition, many modifications may be made to adapt a particular step, structure, or material to the teachings of the invention without departing from its scope. Therefore, it is intended that the invention not be limited to the particular embodiment disclosed, but that the invention will include all embodiments falling within the scope of the appended claims.